Noname manuscript No.
(will be inserted by the editor)

Automatic Quality Assessment for Speech Translation Using Joint ASR and MT Features

Ngoc-Tien Le · Benjamin Lecouteux · Laurent Besacier

Received : date / Accepted : date

arXiv:1609.06049v1 [cs.CL] 20 Sep 2016

Abstract This paper addresses automatic quality assessment of spoken language translation (SLT). This relatively new task is defined and formalized as a sequence labeling problem where each word in the SLT hypothesis is tagged as *good* or *bad* according to a large feature set. We propose several word confidence estimators (WCE) based on our automatic evaluation of transcription (ASR) quality, translation (MT) quality, or both (combined ASR+MT). This research work is possible because we built a specific corpus which contains 6.7k utterances for which a quintuplet containing: ASR output, verbatim transcript, text translation, speech translation and post-edition of translation is built. The conclusion of our multiple experiments using joint ASR and MT features for WCE is that MT features remain the most influent while ASR feature can bring interesting complementary information. Our robust quality estimators for SLT can be used for re-scoring speech translation graphs or for providing feedback to the user in interactive speech translation or computer-assisted speech-to-text scenarios.

Keywords Quality estimation · Word confidence estimation (WCE) · Spoken Language Translation (SLT) · Joint Features · Feature Selection

1 Introduction

Automatic quality assessment of spoken language translation (SLT), also named confidence estimation (CE), is an important topic because it allows to know if a system produces (or not) user-acceptable outputs. In interactive speech to speech translation, CE helps to judge if a translated turn is uncertain

Firstname Lastname
Laboratoire d'Informatique de Grenoble, University of Grenoble Alpes, France
Building IMAG, 700 Centrale, 38401 Saint Martin d'Hères
Tel. : +33 457421454
E-mail: firstname.lastname@imag.fr

(and ask the speaker to rephrase or repeat). For speech-to-text applications, CE may tell us if output translations are worth being corrected or if they require retranslation from scratch. Moreover, an accurate CE can also help to improve SLT itself through a second-pass N-best list re-ranking or search graph re-decoding, as it has already been done for text translation in [2] and [19], or for speech translation in [4]. Consequently, building a method which is capable of pointing out the correct parts as well as detecting the errors in a speech translated output is crucial to tackle above issues.

Given signal x_f in the source language, spoken language translation (SLT) consists in finding the most probable target language sequence $\hat{e} = (e_1, e_2, ..., e_N)$ so that

$$\hat{e} = \underset{e}{\operatorname{argmax}}\{p(e/x_f, f)\} \tag{1}$$

where $f = (f_1, f_2, ..., f_M)$ is the transcription of x_f.

Now, if we perform confidence estimation at the "words" level, the problem is called Word-level Confidence Estimation (WCE) and we can represent this information as a sequence q (same length N of $\hat{e}$) where $q = (q_1, q_2, ..., q_N)$ and $q_i \in \{good, bad\}$ [1].

Then, integrating automatic quality assessment in our SLT process can be done as following:

$$\hat{e} = \underset{e}{\operatorname{argmax}} \sum_{q} p(e, q/x_f, f) \tag{2}$$

$$\hat{e} = \underset{e}{\operatorname{argmax}} \sum_{q} p(q/x_f, f, e) * p(e/x_f, f) \tag{3}$$

$$\hat{e} \approx \underset{e}{\operatorname{argmax}}\{\max_{q}\{p(q/x_f, f, e) * p(e/x_f, f)\}\} \tag{4}$$

In the product of (4), the SLT component $p(e/x_f, f)$ and the WCE component $p(q/x_f, f, e)$ contribute together to find the best translation output $\hat{e}$. In the past, WCE has been treated separately in ASR or MT contexts and we propose here a joint estimation of word confidence for a spoken language translation (SLT) task involving both ASR and MT.

This journal paper is an extended version of a paper published at ASRU 2015 last year [4] but we focus more on the WCE component and on the best approaches to estimate $p(q/x_f, f, e)$ accurately.

Contributions The main contributions of this journal paper are the following:

– A corpus (distributed to the research community [2]) dedicated to WCE for SLT was initially published in [3]. We present, in this paper, its extension from 2643 to 6693 speech utterances.

1. q_i could be also more than 2 labels, or even scores but this paper only deals with error detection (binary set of labels)

2. https://github.com/besacier/WCE-SLT-LIG

- While our previous work on quality assessment was based on two separate WCE classifiers (one for quality assessment in ASR and one for quality assessment in MT), we propose here a unique *joint* model based on different feature types (ASR and MT features).
- This *joint* model allows us to operate feature selection and analyze which features (from ASR or MT) are the most efficient for quality assessment in speech translation.
- We also experiment with two ASR systems that have different performance in order to analyze the behavior of our SLT quality assessment algorithms at different levels of word error rate (WER).

Outline The outline of this paper goes simply as follows: section 2 reviews the state-of-the-art on confidence estimation for ASR and MT. Our WCE system using multiple features is then described in section 3. The experimental setup (notably our specific WCE corpus) is presented in section 4 while section 5 evaluates our joint WCE system. Feature selection for quality assessment in speech translation is analyzed in section 6 and finally, section 7 concludes this work and gives some perspectives.

2 Related work on confidence estimation for ASR and MT

Several previous works tried to propose effective confidence measures in order to detect errors on ASR outputs. Confidence measures are introduced for Out-Of-Vocabulary (OOV) detection by [1]. [27] extends the previous work and introduces the use of word posterior probability (WPP) as a confidence measure for speech recognition. Posterior probability of a word is most of the time computed using the hypothesis word graph [10]. Also, more recent approaches [16] for confidence measure estimation use side-information extracted from the recognizer: normalized likelihoods (WPP), the number of competitors at the end of a word (hypothesis density), decoding process behavior, linguistic features, acoustic features (acoustic stability, duration features) and semantic features.

In parallel, the Workshop on Machine Translation (WMT) introduced in 2013 a WCE task for machine translation. [9] [21] employed the Conditional Random Fields (CRF) [12] model as their machine learning method to address the problem as a sequence labelling task. Meanwhile, [5] extended their initial proposition by dynamic training with adaptive weight updates in their neural network classifier. As far as prediction indicators are concerned, [5] proposed seven word feature types and found among them the "common cover links" (the links that point from the leaf node containing this word to other leaf nodes in the same subtree of the syntactic tree) the most outstanding. [9] focused only on various n-gram combinations of target words. Inheriting most of previously-recognized features, [21] integrated a number of new indicators relying on graph topology, pseudo reference, syntactic behavior (constituent label, distance to the semantic tree root) and polysemy characteristic. The estimation of the confidence score uses mainly classifiers like Conditional Random

Fields [9,18], Support Vector Machines [13] or Perceptron [5]. Some investigations were also conducted to determine which features seem to be the most relevant. [13] proposed to filter features using a forward-backward algorithm to discard linearly correlated features. Using Boosting as learning algorithm, [20] was able to take advantage of the most significant features.

Finally, several toolkits for WCE were recently proposed: *TranscRater* for ASR [3], Marmot for MT [4] as well as WCE-LIG [25] [5] that will be used to extract MT features in the experiments of this journal paper.

To our knowledge, the first attempt to design WCE for speech translation, using both ASR and MT features, is our own work [3,4] which is further extended in this journal paper submission.

3 Building an efficient quality assessment (WCE) system

The WCE component solves the equation:

$$\hat{q} = \underset{q}{\operatorname{argmax}}\{p_{SLT}(q/x_f, f, e)\} \tag{5}$$

where $q = (q_1, q_2, ..., q_N)$ is the sequence of quality labels on the target language. This is a sequence labelling task that can be solved with several machine learning techniques such as Conditional Random Fields (CRF) [12]. However, for that, we need a large amount of training data for which a quadruplet (x_f, f, e, q) is available. In this work, we will use a corpus extended from [3] which contains 6.7k utterances. We will investigate if this amount of data is enough to evaluate and test a joint model $p_{SLT}(q/x_f, f, e)$.

As it is much easier to obtain data containing either the triplet (x_f, f, q) (automatically transcribed speech with manual references and quality labels infered from word error rate estimation) or the triplet (f, e, q) (automatically translated text with manual post-editions and quality labels infered using tools such as TERpA [26]) we can also recast the WCE problem with the following equation:

$$\hat{q} = \underset{q}{\operatorname{argmax}}\{p_{ASR}(q/x_f, f)^{\alpha} * p_{MT}(q/e, f)^{1-\alpha}\} \tag{6}$$

where α is a weight giving more or less importance to WCE_{ASR} (quality assesment on transcription) compared to WCE_{MT} (quality assesment on translation). It is important to note that $p_{ASR}(q/x_f, f)$ corresponds to the quality estimation of the words in the target language based on features calculated on the source language (ASR). For that, what we do is projecting source quality scores to the target using word-alignment information between e and f sequences. This alternative approach (*equation 6*) will be also evaluated in this work.

3. https://github.com/hlt-mt/TranscRater
4. https://github.com/qe-team/marmot
5. https://github.com/besacier/WCE-LIG

Table 1 List of MT features extracted.

1 Proper Name	10 Stop Word	19 WPP max
2 Unknown Stem	11 Word context Alignments	20 Nodes
3 Num. of Word Occ.	12 POS context Alignments	21 Constituent Label
4 Num. of Stem Occ.	13 Stem context Alignments	22 Distance To Root
5 Polysemy Count – Target	14 Longest Target N-gram Length	23 Numeric
6 Backoff Behaviour – Target	15 Longest Source N-gram Length	24 Punctuation
7 Alignment Features	16 WPP Exact	
8 Occur in Google Translate	17 WPP Any	
9 Occur in Bing Translator	18 WPP min	

In both approaches – *joint* ($p_{SLT}(q/x_f, f, e)$) and *combined* ($p_{ASR}(q/x_f, f)$ + $p_{MT}(q/e, f)$) – some features need to be extracted from ASR and MT modules. They are more precisely detailed in next subsections.

3.1 WCE features for speech transcription (ASR)

In this work, we extract several types of features, which come from the ASR graph, from language model scores and from a morphosyntactic analysis. These features are listed below (more details can be found in [3]):

- Acoustic features: word duration (**F-dur**).
- Graph features (extracted from the ASR word confusion networks): number of alternative (**F-alt**) paths between two nodes; word posterior probability (**F-post**).
- Linguistic features (based on probabilities by the language model): word itself (**F-word**), 3-gram probability (**F-3g**), log probability (**F-log**), back-off level of the word (**F-back**), as proposed in [6],
- Lexical Features: Part-Of-Speech (POS) of the word (**F-POS**),
- Context Features: Part-Of-Speech tags in the neighborhood of a given word (**F-context**).

For each word in the ASR hypothesis, we estimate the 9 features (F-Word; F-3g; F-back; F-log; F-alt; F-post; F-dur; F-POS; F-context) previously described.

In a preliminary experiment, we will evaluate these features for quality assessment in ASR only (WCE_{ASR} task). Two different classifiers will be used: a variant of boosting classification algorithm called *bonzaiboost* [14] (implementing the boosting algorithm *Adaboost.MH* over deeper trees) and the Conditional Random Fields [12].

3.2 WCE features for machine translation (MT)

A number of knowledge sources are employed for extracting features, in a total of 24 major feature types, see Table 1.

It is important to note that we extract features regarding *tokens* in the machine translation (MT) hypothesis sentence. In other words, one feature is

extracted for each token in the MT output. So, in the Table 1, *target* refers to the feature coming from the MT hypothesis and *source* refers to a feature extracted from the source word aligned to the considered target word. More details on some of these features are given in the next subsections.

3.2.1 Internal Features

These features are given by the Machine Translation system, which outputs additional data like N-best list.

Word Posterior Probability (WPP) and **Nodes** features are extracted from a confusion network, which comes from the output of the machine translation N-best list. **WPP Exact** is the WPP value for each word concerned at the exact same position in the graph. **WPP Any** extracts the same information at any position in the graph. **WPP Min** gives the smallest WPP value concerned by the transition and **WPP Max** its maximum.

3.2.2 External Features

Below is the list of the external features used:

- **Proper Name**: indicates if a word is a proper name (same binary features are extracted to know if a token is **Numerical**, **Punctuation** or **Stop Word**).
- **Unknown Stem**: informs whether the stem of the considered word is known or not.
- **Number of Word/Stem Occurrences**: counts the occurrences of a word/stem in the sentence.
- **Alignment context features**: these features (#11-13 in Table 1) are based on collocations and proposed by [2]. Collocations could be an indicator for judging if a target word is generated by a particular source word. We also apply the reverse, the collocations regarding the source side (#7 in Table 1 - simply called **Alignment Features**):
  - *Source alignment context features*: the combinations of the target word, the source word (with which it is aligned), and one source word before and one source word after (left and right contexts, respectively).
  - *Target alignment context features*: the combinations of the source word, the target word (with which it is aligned), and one target word before and one target word after.
- **Longest Target (or Source) N-gram Length**: we seek to get the length (n+1) of the longest left sequence (w_{i-n}) concerned by the current word (w_i) and known by the language model (LM) concerned (source and target sides). For example, if the longest left sequence w_{i-2}, w_{i-1}, w_i appears in the target LM, the longest target n-gram value for w_i will be 3. This value ranges from 0 to the max order of the LM concerned. We also extract a redundant feature called **Backoff Behavior Target**.

- The target word's constituent label (**Constituent Label**) and its depth in the constituent tree (**Distance to Root**) are extracted using a syntactic parser.
- **Target Polysemy Count**: we extract the polysemy count, which is the number of meanings of a word in a given language.
- **Occurences in Google Translate** and **Occurences in Bing Translator**: in the translation hypothesis, we (optionally) test the presence of the target word in on-line translations given respectively by *Google Translate* and *Bing Translator*[6].

A very similar feature set was used for a simple WCE_{MT} task (English - Spanish MT, WMT 2013, 2014 quality estimation shared task) and obtained very good performances [17]. This preliminary experience in participating to the WCE shared task in 2013 and 2014 lead us to the following observation: while feature processing is very important to achieve good performance, it requires to call a set of heterogeneous NLP tools (for lexical, syntactic, semantic analyses). Thus, we recently proposed to unify the feature processing, together with the call of machine learning algorithms, in order to facilitate the design of confidence estimation systems. The open-source toolkit proposed (written in *Python* and made available on *github*[7]) integrates some standard as well as in-house features that have proven useful for WCE (based on our experience in WMT 2013 and 2014).

In this paper, we will use only Conditional Random Fields [12] (CRFs) as our machine learning method, with WAPITI toolkit [15], to train our WCE estimator based on MT features.

4 Experimental setup

4.1 Dataset

4.1.1 Starting point: an existing MT Post-edition corpus

For a French-English translation task, we used our SMT system to obtain the translation hypothesis for 10,881 source sentences taken from news corpora of the WMT (Workshop on Machine Translation) evaluation campaign (from 2006 to 2010). Post-editions were obtained from non professional translators using a crowdsourcing platform. More details on the baseline SMT system used can be found in [22] and more details on the post-edited corpus can be found in [23]. It is worth mentionning, however, that a sub-set (311 sentences) of these collected post-editions was assessed by a professional translator and 87.1% of post-editions were judged to improve the hypothesis

6. Using this kind of feature is controversial, however we observed that such features are available in general use case scenarios, so we decided to include them in our experiments. Contrastive results without these 2 features will be also given later on.

7. `http://github.com/besacier/WCE-LIG`

Table 2 Example of training label obtained using TERp-A.

Reference	The	consequence	of	the	fundamentalist	
		S			S	
Hyp After Shift	The	result	of	the	hard-line	
Reference	movement		also	has	its importance	.
	Y	I		D	P	
Hyp After Shift	trend	is	also		important	.

Then, the word label setting for WCE was done using TERp-A toolkit [26]. Table 2 illustrates the labels generated by TERp-A for one hypothesis and post-edition pair. Each word or phrase in the hypothesis is aligned to a word or phrase in the post-edition with different types of edit: "I" (insertions), "S" (substitutions), "T" (stem matches), "Y" (synonym matches), and "P" (phrasal substitutions). The lack of a symbol indicates an exact match and will be replaced by "E" thereafter. We do not consider the words marked with "D" (deletions) since they appear only in the reference. However, later on, we will have to train binary classifiers (*good*/*bad*) so we re-categorize the obtained 6-label set into binary set: The E, T and Y belong to the *good* (G), whereas the S, P and I belong to the *bad* (B) category.

4.1.2 Extending the corpus with speech recordings and transcripts

The *dev* set and *tst* set of this corpus were recorded by french native speakers. Each sentence was uttered by 3 speakers, leading to 2643 and 4050 speech recordings for *dev* set and *tst* set, respectively. For each speech utterance, a quintuplet containing: ASR output (f_{hyp}), verbatim transcript (f_{ref}), English text translation output ($e_{hyp_{mt}}$), speech translation output ($e_{hyp_{slt}}$) and post-edition of translation (e_{ref}), was made available. This corpus is available on a *github* repository [8]. More details are given in table 3. The total length of the *dev* and *tst* speech corpus obtained are 16h52, since some utterances were pretty long.

Table 3 Details on our *dev* and *test* corpora for SLT.

Corpus	**#sentences**	**#speech recordings**	**#speakers**	**Duration**
dev	881	2643	15 (9 women + 6 men)	5h51
tst	1350	4050	27 (11 women + 16 men)	11h01

8. https://github.com/besacier/WCE-SLT-LIG/

Table 4 Details on language models (LM) used in our two ASR systems.

LM	1-gram	2-grams	3-grams
small ($ASR1$)	62K	1M	59M
big ($ASR2$)	95K	49M	301M

4.2 ASR Systems

To obtain the speech transcripts (f_{hyp}), we built a French ASR system based on KALDI toolkit [24]. Acoustic models are trained using several corpora (ESTER, REPERE, ETAPE and BREF120) representing more than 600 hours of french transcribed speech.

The baseline GMM system is based on mel-frequency cepstral coefficient (MFCC) acoustic features (13 coefficients expanded with delta and double delta features and energy : 40 features) with various feature transformations including linear discriminant analysis (LDA), maximum likelihood linear transformation (MLLT), and feature space maximum likelihood linear regression (fMLLR) with speaker adaptive training (SAT). The GMM acoustic model makes initial phoneme alignments of the training data set for the following DNN acoustic model training.

The speech transcription process is carried out in two passes: an automatic transcript is generated with a GMM-HMM model of 43182 states and 250000 Gaussians. Then word graphs outputs obtained during the first pass are used to compute a fMLLR-SAT transform on each speaker. The second pass is performed using DNN acoustic model trained on acoustic features normalized with the fMLLR matrix.

CD-DNN-HMM acoustic models are trained (43 182 context-dependent states) using GMM-HMM topology.

We propose to use two 3-gram language models trained on French ESTER corpus [8] as well as on French Gigaword (vocabulary size are respectively 62k and 95k). The ASR systems LM weight parameters are tuned through WER on the *dev* corpus. Details on these two language models can be found in table 4.

In our experiments we propose two ASR systems based on the previously described language models. The first system ($ASR1$) uses the small language model allowing a fast ASR system (about 2x Real Time), while in the second system lattices are rescored with a big language model (about 10x Real Time) during a third pass.

Table 5 presents the performances obtained by two above ASR systems.

Table 5 ASR performance (WER) on our *dev* and *test* set for the two different ASR systems

Task	*dev* set	*tst* set
ASR1	21.86%	17.37%
ASR2	16.90%	12.50%

These WER may appear as rather high according to the task (transcribing read news). A deeper analysis shows that these news contain a lot of foreign named entities, especially in our *dev* set. This part of the data is extracted from French medias dealing with european economy in EU. This could also explain why the scores are significantly different between *dev* and *test* sets. In addition, automatic post-processing is applied to ASR output in order to match requirements of standard input for machine translation.

4.3 SMT System

We used *moses* phrase-based translation toolkit [11] to translate French ASR into English (e_{hyp}). This medium-size system was trained using a subset of data provided for IWSLT 2012 evaluation [7]: Europarl, Ted and News-Commentary corpora. The total amount is about 60M words. We used an adapted target language model trained on specific data (News Crawled corpora) similar to our evaluation corpus (see [22]). This standard SMT system will be used in all experiments reported in this paper.

4.4 Obtaining quality assessment labels for SLT

After building an ASR system, we have a new element of our desired quintuplet: the ASR output f_{hyp}. It is the noisy version of our already available verbatim transcripts called f_{ref}. This ASR output (f_{hyp}) is then translated by the exact same SMT system [22] already mentionned in subsection 4.3. This new output translation is called $e_{hyp_{slt}}$ and it is a degraded version of $e_{hyp_{mt}}$ (translation of f_{ref}).

At this point, a strong assumption we made has to be revealed: we re-used the post-editions obtained from the text translation task (called e_{ref}), to infer the quality (G, B) labels of our speech translation output $e_{hyp_{slt}}$. The word label setting for WCE is also done using TERp-A toolkit [26] between $e_{hyp_{slt}}$ and e_{ref}. This assumption, and the fact that initial MT post-edition can be also used to infer labels of a SLT task, is reasonnable regarding results (later presented in table 8 and table 9) where it is shown that there is not a huge difference between the MT and SLT performance (evaluated with BLEU).

The remark above is important and this is what makes the value of this corpus. For instance, other corpora such as the TED corpus compiled by LIUM [9] contain also a quintuplet with ASR output, verbatim transcript, MT output, SLT output and target translation. But there are 2 main differences: first, the target translation is a manual translation of the prior subtitles so this is not a post-edition of an automatic translation (and we have no guarantee that the *good/bad* labels extracted from this will be reliable for WCE training and

9. http://www-lium.univ-lemans.fr/fr/content/corpus-ted-lium

testing); secondly, in our corpus, each sentence is uttered by 3 different speakers which introduces speaker variability in the database and allows us to deal with different ASR outputs for a single source sentence.

4.5 Final corpus statistics

The final corpus obtained is summarized in table 6, where we also clarify how the WCE labels were obtained. For the test set, we now have all the data needed to evaluate WCE for 3 tasks:

- **ASR**: extract *good/bad* labels by calculating WER between f_{hyp} and f_{ref},
- **MT**: extract *good/bad* labels by calculating TERp-A between $e_{hyp_{mt}}$ and e_{ref},
- **SLT**: extract *good/bad* labels by calculating TERp-A between $e_{hyp_{slt}}$ and e_{ref}.

Table 6 Overview of our post-edition corpus for SLT.

Data	**# *dev* utt**	**# *test* utt**	**method to obtain WCE labels**
f_{ref}	881	1350	
f_{hyp}	881*3	1350*3	wer(f_{hyp}, f_{ref})
$e_{hyp_{mt}}$	881	1350	terpa($e_{hyp_{mt}}$, e_{ref})
$e_{hyp_{slt}}$	881*3	1350*3	terpa($e_{hyp_{slt}}$, e_{ref})
e_{ref}	881	1350	

Table 7 gives an example of the quintuplet available in our corpus. One transcript (f_{hyp1}) has 1 error while the other one (f_{hyp2}) has 4. This leads to respectively 2 B labels ($e_{hyp_{slt1}}$) and 4 B labels ($e_{hyp_{slt2}}$) in the speech translation output, while $e_{hyp_{mt}}$ has only one B label.

Table 7 Example of quintuplet with associated labels.

f_{ref}	quand	notre	cerveau	chauffe
f_{hyp1}	*comme*	notre	cerveau	chauffe
labels ASR	B	G	G	G
f_{hyp2}	*qu'*	*entre*	*serbes*	*au* chauffe
labels ASR	B	B	B	B G
$e_{hyp_{mt}}$	when	our	brains	*chauffe*
labels MT	G	G	G	B
$e_{hyp_{slt1}}$	*as*	our	brains	*chauffe*
labels SLT	B	G	G	B
$e_{hyp_{slt2}}$	*between*	*serbs*	*in*	*chauffe*
labels SLT	B	B	B	B
e_{ref}	when	our	brain	heats up

Table 8 and table 9 summarize baseline ASR, MT and SLT performances obtained on our corpora, as well as the distribution of good (G) and bad (B) labels inferred for both tasks. Logically, the percentage of (B) labels increases from MT to SLT task in the same conditions.

Table 8 MT and SLT performances on our *dev* set.

Task	**ASR (WER)**	**MT (BLEU)**	**% G (good)**	**% B (bad)**
MT	0%	49.13%	76.93%	23.07%
SLT (ASR1)	21.86%	26.73%	62.03%	37.97%
SLT (ASR2)	16.90%	28.89%	63.87%	36.13%

Table 9 MT and SLT performances on our *tst* set.

Task	**ASR (WER)**	**MT (BLEU)**	**% G (good)**	**% B (bad)**
MT	0%	57.87%	81.58%	18.42%
SLT ($ASR1$)	17.37%	30.89%	61.12%	38.88%
SLT ($ASR2$)	12.50%	33.14%	62.77%	37.23%

5 Experiments on WCE for SLT

5.1 SLT quality assessment using only MT or ASR features

We first report in Table 10 the baseline WCE results obtained using MT or ASR features separately. In short, we evaluate the performance of 4 WCE systems for different tasks:

- The first and second systems (WCE for ASR / ASR feat.) use ASR features described in section 3.1 with two different classifiers (CRF or Boosting).
- The third system (WCE for SLT / MT feat.) uses only MT features described in section 3.2 with CRF classifier.
- The fourth system (WCE for SLT / ASR feat.) uses only ASR features described in section 3.1 with CRF classifier (so this is predicting SLT output confidence using only ASR confidence features!). Word alignment information between f_{hyp} and e_{hyp} is used to project the WCE scores coming from ASR, to the SLT output,

In all experiments reported in this paper, we evaluate the performance of our classifiers by using the average between the F-measure for *good* labels and the F-measure for *bad* labels that are calculated by the common evaluation metrics: Precision, Recall and F-measure for *good*/*bad* labels. Since two ASR

systems are available, *F-mes1* is obtained for SLT based on $ASR1$ whereas *F-mes2* is obtained for SLT based on $ASR2$. For the results of Table 10, the classifier is evaluated on the *tst* part of our corpus and trained on the *dev* part.

Table 10 WCE performance with different feature sets for *tst* set (training is made on *dev* set) - *for MT feat, removing *OccurInGoogleTranslate* and *OccurInBingTranslate* features lead to 59.40% and 58.11% for *F-mes1* and *F-mes2* respectively

task feat. type	**WCE for ASR** ASR feat. $p(q/x_f, f)$ (CRFs)	**WCE for ASR** ASR feat. $p(q/x_f, f)$ (Boosting)	**WCE for SLT** MT feat. $p(q/f, e)$	**WCE for SLT** ASR feat. $p(q/x_f, f)$ projected to e
F-mes1	68.71%	64.27%	60.55%*	49.67%
F-mes2	59.83%	62.61%	59.83%*	44.56%

Concerning WCE for ASR, we observe that Fmeasure decreases when ASR WER is lower (*F-mes2*<*F-mes1* while $WER_{ASR2} < WER_{ASR1}$). So quality assessment in ASR seems to become harder as the ASR system improves. This could be due to the fact that the ASR1 errors recovered by bigger LM in ASR2 system were easier to detect. Anyway, this conclusion should be considered with caution since both results (*F-mes1* and *F-mes2*) are not directly comparable because they are evaluated on different references (proportion of *good*/*bad* labels differ as ASR system differ). The effect of the classifier (CRF or Boosting) is not conclusive since CRF is better for *F-mes1* and worse for *F-mes2*. Anyway, we decide to use CRF for all our future experiments since this is the classifier integrated in WCE-LIG [25] toolkit.

Concerning WCE for SLT, we observe that Fmeasure is better using MT features rather than ASR features (quality assessment for SLT more dependent of MT features than ASR features). Again, Fmeasure decreases when ASR WER is lower (*F-mes2*<*F-mes1* while $WER_{ASR2} < WER_{ASR1}$). For MT features, removing *OccurInGoogleTranslate* and *OccurInBingTranslate* features lead to 59.40% and 58.11% for *F-mes1* and *F-mes2* respectively.

In the next subsection, we try to see if the use of both MT and ASR features improves quality assessment for SLT.

5.2 SLT quality assessment using both MT and ASR features

We now report in Table 12 WCE for SLT results obtained using both MT and ASR features. More precisely we evaluate two different approaches (*combination* and *joint*):

- The first system (WCE for SLT / MT+ASR feat.) combines the output of two separate classifiers based on ASR and MT features. In this approach, ASR-based confidence score of the source is projected to the target SLT output and combined with the MT-based confidence score as shown in *equation 6* (we did not tune the α coefficient and set it *a priori* to 0.5).

– The second system (joint feat.) trains a single WCE system for SLT (evaluating $p(q/x_f, f, e)$ as in *equation 5* using joint ASR features and MT features. All ASR features are projected to the target words using automatic word alignments. However, a problem occur when a target word does not have any source word aligned to it. In this case, we decide to duplicate the ASR features of its previous target word. Another problem occur when a target word is aligned to more than one source word. In that case, there are several strategies to infer the 9 ASR features: average or max over numerical values, selection or concatenation over symbolic values (for F-word and F-POS), etc. Three different variants of these strategies (shown in Table 11) are evaluated here.

Table 11 Different strategies to project ASR features to a target word when it is aligned to more than one source word. *It should be noted that **F-context** features are the combinations of the source word (**F-word**) and one POS of source word (**F-POS**) before and one POS of source word (**F-POS**) after.

***ASR* Feat**	**Joint 1**	**Joint 2**	**Joint 3**
F-post	avg(F-post1, F-post2)	avg(F-post1, F-post2)	avg(F-post1, F-post2)
F-log	avg(F-log1, F-log2)	avg(F-log1, F-log2)	avg(F-log1, F-log2)
F-back	avg(F-back1, F-back2)	avg(F-back1, F-back2)	avg(F-back1, F-back2)
F-dur	max(F-dur1, F-dur2)	max(F-dur1, F-dur2)	max(F-dur1, F-dur2)
F-3g	max(F-3g1, F-3g2)	max(F-3g1, F-3g2)	max(F-3g1, F-3g2)
F-alt	max(F-alt1, F-alt2)	max(F-alt1, F-alt2)	max(F-alt1, F-alt2)
F-word	F-word1	F-word2	F-word1_F-word2
F-POS	F-POS1	F-POS2	F-POS1_F-POS2
F-context	F-context*	F-context*	F-context*

Table 12 WCE performance with combination (MT+ASR) or joint (MT,ASR) feature sets for *tst* set (training is made on *dev* set) - * For *Joint 1* feat, removing *OccurInGoogleTranslate* and *OccurInBingTranslate* features lead to 59.14% and 57.75% for *F-mes1* and *F-mes2* respectively.

task feat. type	**WCE for SLT** MT+ASR feat. $p_{ASR}(q/x_f, f)^{\alpha}$ $*p_{MT}(q/e, f)^{1-\alpha}$	**WCE for SLT** Joint feat. 1 $p(q/x_f, f, e)$	**WCE for SLT** Joint feat. 2 $p(q/x_f, f, e)$	**WCE for SLT** Joint feat. 3 $p(q/x_f, f, e)$
F-mes1	52.99%	**60.29%***	60.17%	60.23%
F-mes2	48.46%	**59.23%***	59.20%	58.99%

The results of Table 12 show that joint ASR and MT features do not improve WCE performance: *F-mes1* and *F-mes2* are slightly worse than those of table 9 (WCE for SLT / MT features only). We also observe that simple combination (MT+ASR) degrades the WCE performance. This latter observation may be due to different behaviors of WCE_{MT} and WCE_{ASR} classifiers which makes the weighted combination ineffective. Moreover, the disappoint-

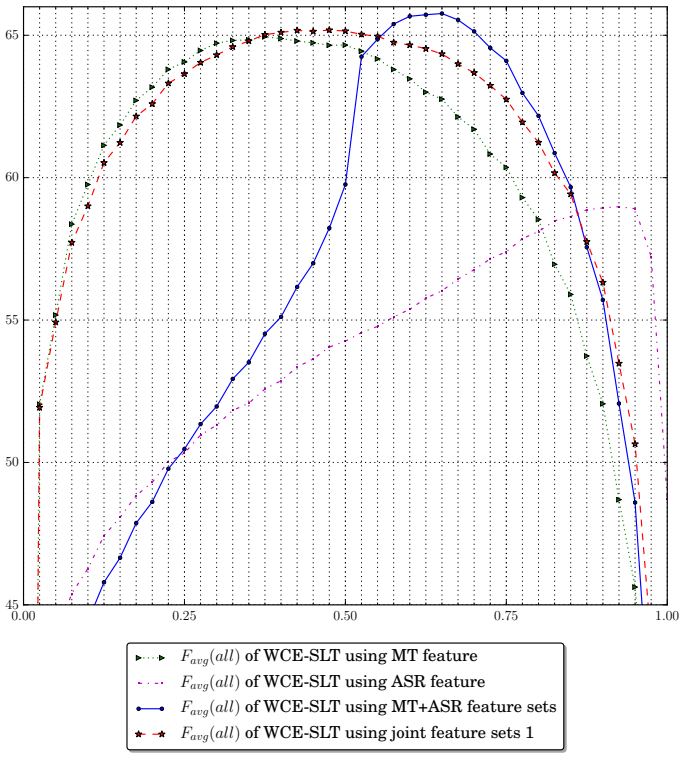

Fig. 1 Evolution of system performance (y-axis - *F-mes1* - ASR1) for *dev* corpus (2683 utt) along decision threshold variation (x-axis) - training is made on 4050 utt.

ing performance of our joint classifier may be due to an insufficient training set (only 2683 utterances in *dev*!). Finally, removing *OccurInGoogleTranslate* and *OccurInBingTranslate* features for *Joint* lowered *F-mes* between 1% and 1.5%.

These observations lead us to investigate the behaviour of our WCE approaches for a large range of *good/bad* decision threshold and with a new protocol where we reverse *dev* and *tst*. So, in the next experiments of this subsection, we will report WCE evaluation results obtained on *dev* (2683 utt.) with classifiers trained on *tst* (4050 utt.). Finally, the different strategies used to project ASR features when a target word is aligned to more than one source word do not lead to very different performance: we will use strategy *joint 1* in the future.

While the previous tables provided WCE performance for a single point of interest (*good/bad* decision threshold set to 0.5), the curves of figures 1 and 2 show the full picture of our WCE systems (for SLT) using speech transcriptions systems $ASR1$ and $ASR2$, respectively. We observe that the classifier based on ASR features has a very different behaviour than the classifier based on MT features which explains why their simple combination (MT+ASR) does not

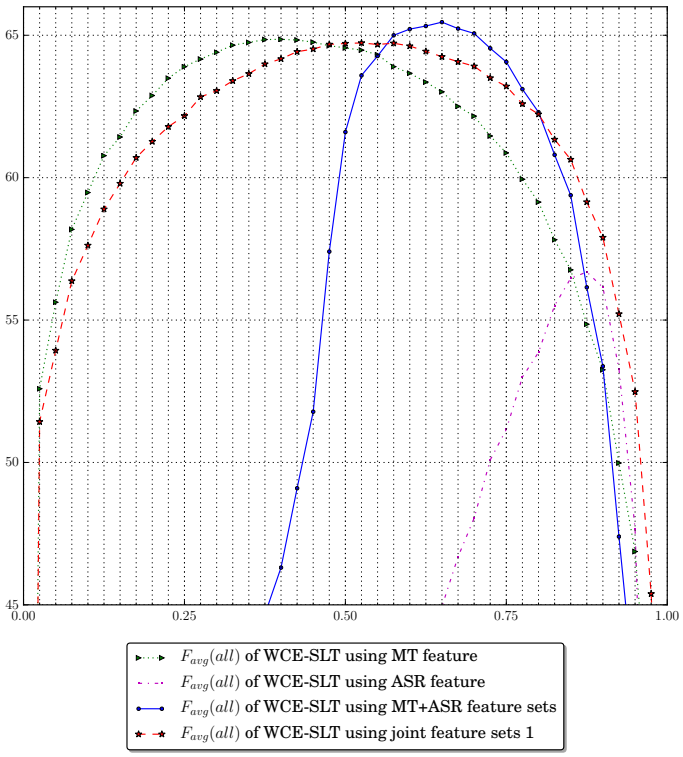

Fig. 2 Evolution of system performance (y-axis - *F-mes2* - ASR2) for *dev* corpus (2683 utt) along decision threshold variation (x-axis) - training is made on 4050 utt.

work very well for the default decision threshold (0.5). However, for threshold above 0.5, the use of both ASR and MT features is beneficial. This is interesting because higher thresholds improves the Fmeasure on *bad* labels (so improves error detection). Both curves are similar whatever the ASR system used. These results suggest that with enough development data for appropriate threshold tuning (which we do not have for this very new task), the use of both ASR and MT features should improve error detection in speech translation (blue and red curves are above the green curve for higher decision threshold [10]). We also analyzed the *Fmeasure* curves for *bad* and *good* labels separately [11]: if we consider, for instance *ASR*1 system, for decision threshold equals to 0.75, the Fmeasure on *bad* labels is equivalent (60%) for 3 systems (*Joint, MT+ASR* and *MT*) while the Fmeasure on *good* labels is 61% when using *MT* features only, 66% when using *Joint* features and 68% when using *MT+ASR* features. In other words, for a fixed performance on *bad* labels, the Fmeasure on *good* labels is improved using all information available (ASR and MT features).

10. Corresponding to optimization of the Fmeasure on *bad* labels (errors)
11. Not reported here due to space constraints.

Finally, if we focus on *Joint* versus *MT+ASR*, we notice that the range of the threshold where performance are stable is larger for *Joint* than for *MT+ASR*.

6 Feature Selection

In this section, we try to better understand the contribution of each (ASR or MT) feature by applying feature selection on our joint WCE classifier. In these experiments, we decide to keep *OccurInGoogleTranslate* and *OccurInBingTranslate* features.

We choose the Sequential Backward Selection (SBS) algorithm which is a top-down algorithm starting from a feature set noted Y_k (which denotes the set of all features) and sequentially removing the most irrelevant one (x) that maximizes the Mean F-Measure, $MF(Y_k - x)$. In our work, we examine until the set Y_k contains only one remaining feature. Algorithm 1 summarizes the whole process.

Algorithm 1 Sequential Backward Selection (SBS) algorithm for feature selection. Y_k denotes the set of all features and x is the feature removed at each step of the algorithm

```
while size of Y_k > 0 do
    maxval = 0
    for x ∈ Y_k do
        if maxval < MF(Y_k − x) then
            maxval ← MF(Y_k − x)
            worstfeat ← x
        end if
    end for
    remove worstfeat from Y_k
end while
```

The results of the SBS algorithm can be found in table 13 which ranks all joint features used in WCE for SLT by order of importance after applying the algorithm on *dev*. We can see that the SBS algorithm is not very stable and is clearly influenced by the ASR system ($ASR1$ or $ASR2$) considered in SLT. Anyway, if we focus on the features that are in the top-10 best in both cases, we find that the most relevant ones are:

- *Occur in Google Translate* and *Occur in Bing Translate* (diagnostic from other MT systems),
- *Longest Source N-gram Length, Target Backoff Behaviour* (source or target N-gram features)
- *Stem Context Alignment* (source-target alignment feature)

We also observe that the most relevant ASR features (in bold in table 13) are *F-3g, F-POS* and *F-back* (lexical and linguistic features) whereas ASR acoustic and graph based features are among the worst (*F-post, F-alt, F-dur*). So, in our experimental setting, it seems that MT features are more influent than ASR features. Another surprising result is the relatively low rank of word

Table 13 Rank of each feature according to the Sequential Backward Selection algorithm - WCE for SLT task - Joint (ASR,MT) features used - Feature selection applied on *dev* corpus for both *ASR1* and *ASR2* - ASR features are in bold.

Rank *ASR1*	**Rank *ASR2***	**Feature**	**Rank *ASR1***	**Rank *ASR2***	**Feature**
1	4	Occur in Google Translate	18	14	Numeric
2	2	Longest Source N-gram Length	19	30	Proper Name
3	5	Target Backoff Behaviour	20	20	Unknown Stem
4	22	Constituent Label	21	24	Number of Word Occurrences
5	1	Occur in Bing Translate	22	23	**F-alt**
6	11	**F-3g**	23	15	Nodes
7	16	WPP Exact	24	8	**F-log**
8	7	Stem Context Alignment	25	32	**F-context**
9	17	WPP Max	26	19	Longest Target N-gram Length
10	12	Number of Stem Occurrences	27	6	WPP Any
11	21	Polysemy Count - Target	28	29	POS Context Alignment
12	3	**F-POS**	29	10	**F-post**
13	18	Stop Word	30	28	Word Context Alignment
14	25	Distance to Root	31	31	**F-dur**
15	13	**F-back**	32	9	Alignment Features
16	26	WPP Min	33	33	**F-word**
17	27	Punctuation			

posterior probability (WPP) features whereas we were expecting to see them among the top features (as shown in [20] where *WPP Any* is among the best features for WCE in MT).

Figure 3 and Figure 4 present the evolution of WCE performance for *dev* and *tst* corpora when feature selection using SBS algorithm is made on *dev*, for $ASR1$ and $ASR2$ systems, respectively. In other words, for these two figures, we apply our SBS algorithm on *dev* which means that feature selection is done on *dev* with classifiers trained on *tst*. After that, the best feature subsets (using 33, 32, 31 until 1 feature only) are applied on *tst* corpus (with classifiers trained on *dev*) [12].

On both figures, we observe that half of the features only contribute to the WCE process since best performances are observed with 10 to 15 features only. We also notice that optimal WCE performance is not necessarily obtained with the full feature set but it can be obtained with a subset of it.

7 Conclusion

7.1 Main contributions

In this paper, we introduced a new quality assessment task: word confidence estimation (WCE) for spoken language translation (SLT). A specific

12. 3 data sets would have been needed to (a) train classifiers, (b) apply feature selection, (c) evaluate WCE performance. Since we only have a *dev* and a *tst* set, we found this procedure acceptable

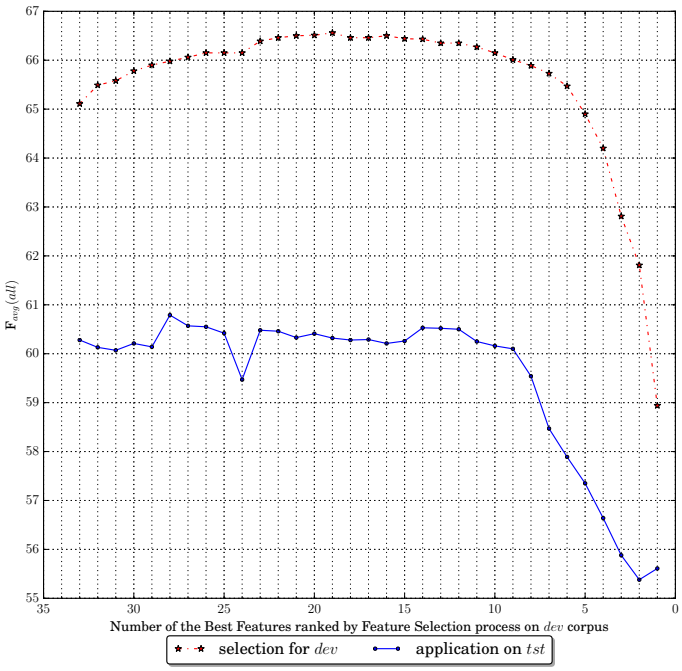

Fig. 3 Evolution of WCE performance for *dev* (features selected) and *tst* corpora when feature selection using SBS algorithm is made on *dev* ($ASR1$ system).

corpus, distributed to the research community [13] was built for this purpose. We formalized WCE for SLT and proposed several approaches based on several types of features: machine translation (MT) based features, automatic speech recognition (ASR) based features, as well as combined or joint features using ASR and MT information. The proposition of a unique *joint* classifier based on different feature types (ASR and MT features) allowed us to operate feature selection and analyze which features (from ASR or MT) are the most efficient for quality assessment in speech translation. Our conclusion is that MT features remain the most influential while ASR feature can bring interesting complementary information. In all our experiments, we systematically evaluated with two ASR systems that have different performance in order to analyze the behavior of our quality assessment algorithms at different levels of word error rate (WER). This allowed us to observe that WCE performance decreases as ASR system improves. For reproducible research, most features [14] and algorithms used in this paper are available through our toolkit called WCE-LIG. This package is made available on a *GitHub* repository [15] under the licence GPL V3. We hope that the availability of our corpus and toolkit could lead, in a near future, to a new shared task dedicated to quality estima-

13. `https://github.com/besacier/WCE-SLT-LIG`
14. MT features already available, ASR features available soon
15. `https://github.com/besacier/WCE-LIG`

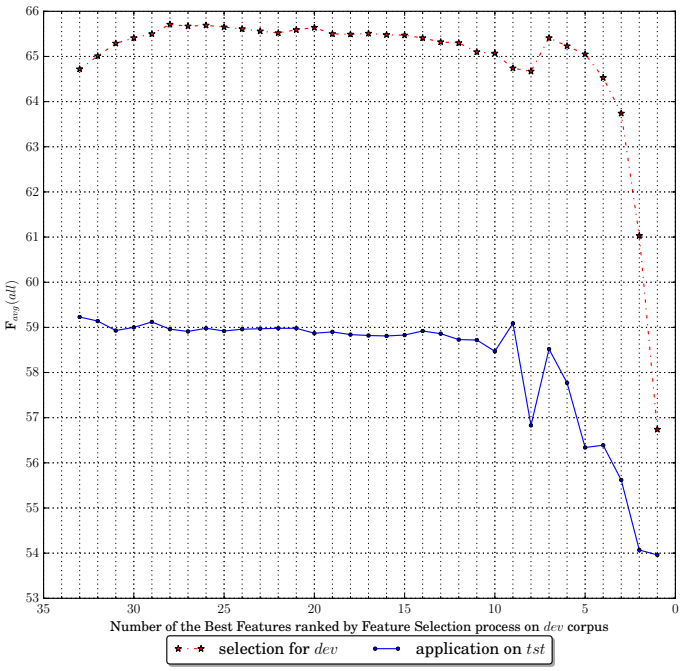

Fig. 4 Evolution of WCE performance for *dev* (features selected) and *tst* corpora when feature selection using SBS algorithm is made on *dev* ($ASR2$ system).

tion for speech translation. Such a shared task could be proposed in avenues such as IWSLT (International Workshop on Spoken Language Translation) or WMT (Workshop on Machine Translation) for instance.

7.2 SLT redecoding using WCE

A direct application of this work is the use of WCE labels to re-decode speech translation graphs and (hopefully) improve speech translation performance. Preliminary results were already obtained and recently published by the authors of this paper [4]. The main idea is to carry a second speech translation pass by considering every word and its quality assessment label, as shown in *equation 4*. The speech translation graph is redecoded following the following principle: words labeled as *good* in the search graph should be "rewarded" by reducing their cost; on the contrary, those labeled as *bad* should be "penalized". To illustrate this direct application of our work, we present examples of speech translation hypotheses (SLT) obtained with or without graph re-decoding in table 14 (table taken from [4]).

Example 1 illustrates a first case where re-decoding allows slightly improving the translation hypothesis. Analysis of the labels from the confidence estimator indicates that the words *a* (start of sentence) and *penalty* were labeled as *bad* here. Thus, a better hypothesis arised from the second pass, although

Table 14 Examples of French SLT hyp with and w/o graph re-decoding - table taken from [4].

f_{ref_1}	une démobilisation des employés peut déboucher sur une démoralisation **mortifère**
f_{hyp_1}	une démobilisation des employés peut déboucher sur une démoralisation **mort y faire**
e_{hyp_1} baseline	a **demobilisation employees** can lead to a **penalty demoralisation**
e_{hyp_1} with re-decoding	a **demobilisation of employees** can lead to a **demoralization death**
e_{ref_1}	**demobilization of employees** can lead to a **deadly demoralization**
f_{ref_2}	celui-ci a indiqué que l'intervention **s'**était parfaitement bien **déroulée** et que les examens post-**opératoires** étaient normaux
f_{hyp_2}	celui-ci a indiqué que l' intervention **c'**était parfaitement bien **déroulés** , et que les examens post **opératoire** étaient normaux.
e_{hyp_2} baseline	**it** has indicated that the speech **that was well** conducted , and that the tests were **normal post route**
e_{hyp_2} with re-decoding	**he** indicated that the intervention **is very well done** , and that the tests **after operating were normal**
e_{ref_2}	**he** indicated that the operation **went perfectly well** and the **post-operative tests were normal**
f_{ref_3}	general motors repousse jusqu'en janvier le plan pour **opel**
f_{hyp_3}	general motors repousse jusqu' en janvier le plan pour **open**
e_{hyp_3} baseline	general motors postponed until january **the plan to open**
e_{hyp_3} with re-decoding	general motors puts until january **terms to open**
e_{ref_3}	general motors postponed until january **the plan for opel**

the transcription error could not be recovered. In example 2, the confidence estimator labeled as *bad* the following word sequences: *it has*, *speech that was* and *post route*. Better translation hypothesis is found after re-decoding (correct pronoun, better quality at the end of sentence). Finally, example 3 shows a case where, this time, the end of the first pass translation deteriorated after re-decoding. Analysis of confidence estimator output shows that the phrase *to open* was (correctly) labeled as *bad*, but the re-decoding gave rise to an even worse hypothesis. The reason is that the system could not recover the named entity *opel* since this word was not in the speech translation graph.

7.3 Other perspectives

In addition to re-decode SLT graphs, our quality assessment system can be used in interactive speech translation scenarios such as news or lectures subtitling, to improve human translator productivity by giving him/her feedback on automatic transcription and translation quality. Another application would be the adaptation of our WCE system to interactive speech-to-speech translation scenarios where feedback on transcription and translation modules

is needed to improve communication. On these latter subjects, it would also be nice to move from a binary (*good* or *bad* labels) to a 3-class decision problem (*good, asr-error, mt-error*). The outcome material of this paper (corpus, toolkit) can be definitely used to address such a new problem.

References

1. Asadi, A., Schwartz, R., Makhoul, J.: Automatic detection of new words in a large vocabulary continuous speech recognition system. Proc. of International Conference on Acoustics, Speech and Signal Processing (1990)
2. Bach, N., Huang, F., Al-Onaizan, Y.: Goodness: A method for measuring machine translation confidence. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics, pp. 211–219. Portland, Oregon (2011)
3. Besacier, L., Lecouteux, B., Luong, N.Q., Hour, K., Hadjsalah, M.: Word confidence estimation for speech translation. In: Proceedings of The International Workshop on Spoken Language Translation (IWSLT). Lake Tahoe, USA (2014)
4. Besacier, L., Lecouteux, B., Luong, N.Q., Le, N.T.: Spoken language translation graphs re-decoding using automatic quality assessment. In: IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU). Scotsdale, Arizona, United States (2015). DOI 10.1109/ASRU.2015.7404804. URL `https://hal.archives-ouvertes.fr/hal-01289158`
5. Bicici, E.: Referential translation machines for quality estimation. In: Proceedings of the Eighth Workshop on Statistical Machine Translation, pp. 343–351. Association for Computational Linguistics, Sofia, Bulgaria (2013). URL `http://www.aclweb.org/anthology/W13-2242`
6. Fayolle, J., Moreau, F., Raymond, C., Gravier, G., Gros, P.: Crf-based combination of contextual features to improve a posteriori word-level confidence measures. In: Interspeech (2010)
7. Federico, M., Cettolo, M., Bentivogli, L., Paul, M., Stüker, S.: Overview of the IWSLT 2012 evaluation campaign. In: In proceedings of the 9th International Workshop on Spoken Language Translation (IWSLT) (2012)
8. Galliano, S., Geoffrois, E., Gravier, G., Bonastre, J.F., Mostefa, D., Choukri, K.: Corpus description of the ester evaluation campaign for the rich transcription of french broadcast news. In: In Proceedings of the 5th international Conference on Language Resources and Evaluation (LREC 2006), pp. 315–320 (2006)
9. Han, A.L.F., Lu, Y., Wong, D.F., Chao, L.S., He, L., Xing, J.: Quality estimation for machine translation using the joint method of evaluation criteria and statistical modeling. In: Proceedings of the Eighth Workshop on Statistical Machine Translation, pp. 365–372. Association for Computational Linguistics, Sofia, Bulgaria (2013). URL `http://www.aclweb.org/anthology/W13-2245`
10. Kemp, T., Schaaf, T.: Estimating confidence using word lattices. Proc. of European Conference on Speech Communication Technology pp. 827–830 (1997)
11. Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., Herbst, E.: Moses: Open source toolkit for statistical machine translation. In: Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, pp. 177–180. Prague, Czech Republic (2007)
12. Lafferty, J., McCallum, A., Pereira, F.: Conditional random fields: Probabilistic models for segmenting et labeling sequence data. In: Proceedings of ICML-01, pp. 282–289 (2001)
13. Langlois, D., Raybaud, S., Smaïli, K.: Loria system for the wmt12 quality estimation shared task. In: Proceedings of the Seventh Workshop on Statistical Machine Translation, pp. 114–119. Baltimore, Maryland USA (2012)

14. Laurent, A., Camelin, N., Raymond, C.: Boosting bonsai trees for efficient features combination : application to speaker role identification. In: Interspeech (2014)
15. Lavergne, T., Cappé, O., Yvon, F.: Practical very large scale crfs. In: Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, pp. 504–513 (2010)
16. Lecouteux, B., Linarès, G., Favre, B.: Combined low level and high level features for out-of-vocabulary word detection. INTERSPEECH (2009)
17. Luong, N.Q., Besacier, L., Lecouteux, B.: Word confidence estimation and its integration in sentence quality estimation for machine translation. In: Proceedings of The Fifth International Conference on Knowledge and Systems Engineering (KSE 2013). Hanoi, Vietnam (2013)
18. Luong, N.Q., Besacier, L., Lecouteux, B.: LIG System for Word Level QE task at WMT14. In: Proceedings of the Ninth Workshop on Statistical Machine Translation, pp. 335–341. Baltimore, Maryland USA (2014)
19. Luong, N.Q., Besacier, L., Lecouteux, B.: Word Confidence Estimation for SMT N-best List Re-ranking. In: Proceedings of the Workshop on Humans and Computer-assisted Translation (HaCaT) during EACL. Gothenburg, Suède (2014). URL `http://hal.inria.fr/hal-00953719`
20. Luong, N.Q., Besacier, L., Lecouteux, B.: Towards accurate predictors of word quality for machine translation: Lessons learned on french - english and english - spanish systems. Data and Knowledge Engineering p. 11 (2015)
21. Luong, N.Q., Lecouteux, B., Besacier, L.: LIG system for WMT13 QE task: Investigating the usefulness of features in word confidence estimation for MT. In: Proceedings of the Eighth Workshop on Statistical Machine Translation, pp. 396–391. Association for Computational Linguistics, Sofia, Bulgaria (2013)
22. Potet, M., Besacier, L., Blanchon, H.: The lig machine translation system for wmt 2010. In: A. Workshop (ed.) Proceedings of the joint fifth Workshop on Statistical Machine Translation and Metrics MATR (WMT2010). Uppsala, Sweden (2010)
23. Potet, M., Emmanuelle E, R., Besacier, L., Blanchon, H.: Collection of a large database of french-english smt output corrections. In: Proceedings of the eighth international conference on Language Resources and Evaluation (LREC). Istanbul, Turkey (2012)
24. Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., Hannemann, M., Motlicek, P., Qian, Y., Schwarz, P., Silovsky, J., Stemmer, G., Vesely, K.: The kaldi speech recognition toolkit. In: IEEE 2011 Workshop on Automatic Speech Recognition and Understanding. IEEE Signal Processing Society (2011). IEEE Catalog No.: CFP11SRW-USB
25. Servan, C., Le, N.T., Luong, N.Q., Lecouteux, B., Besacier, L.: An Open Source Toolkit for Word-level Confidence Estimation in Machine Translation. In: The 12th International Workshop on Spoken Language Translation (IWSLT'15). Da Nang, Vietnam (2015). URL `https://hal.archives-ouvertes.fr/hal-01244477`
26. Snover, M., Madnani, N., Dorr, B., Schwartz, R.: Terp system description. In: MetricsMATR workshop at AMTA (2008)
27. Young, S.R.: Recognition confidence measures: Detection of misrecognitions and out-of-vocabulary words. Proc. of International Conference on Acoustics, Speech and Signal Processing pp. 21–24 (1994)